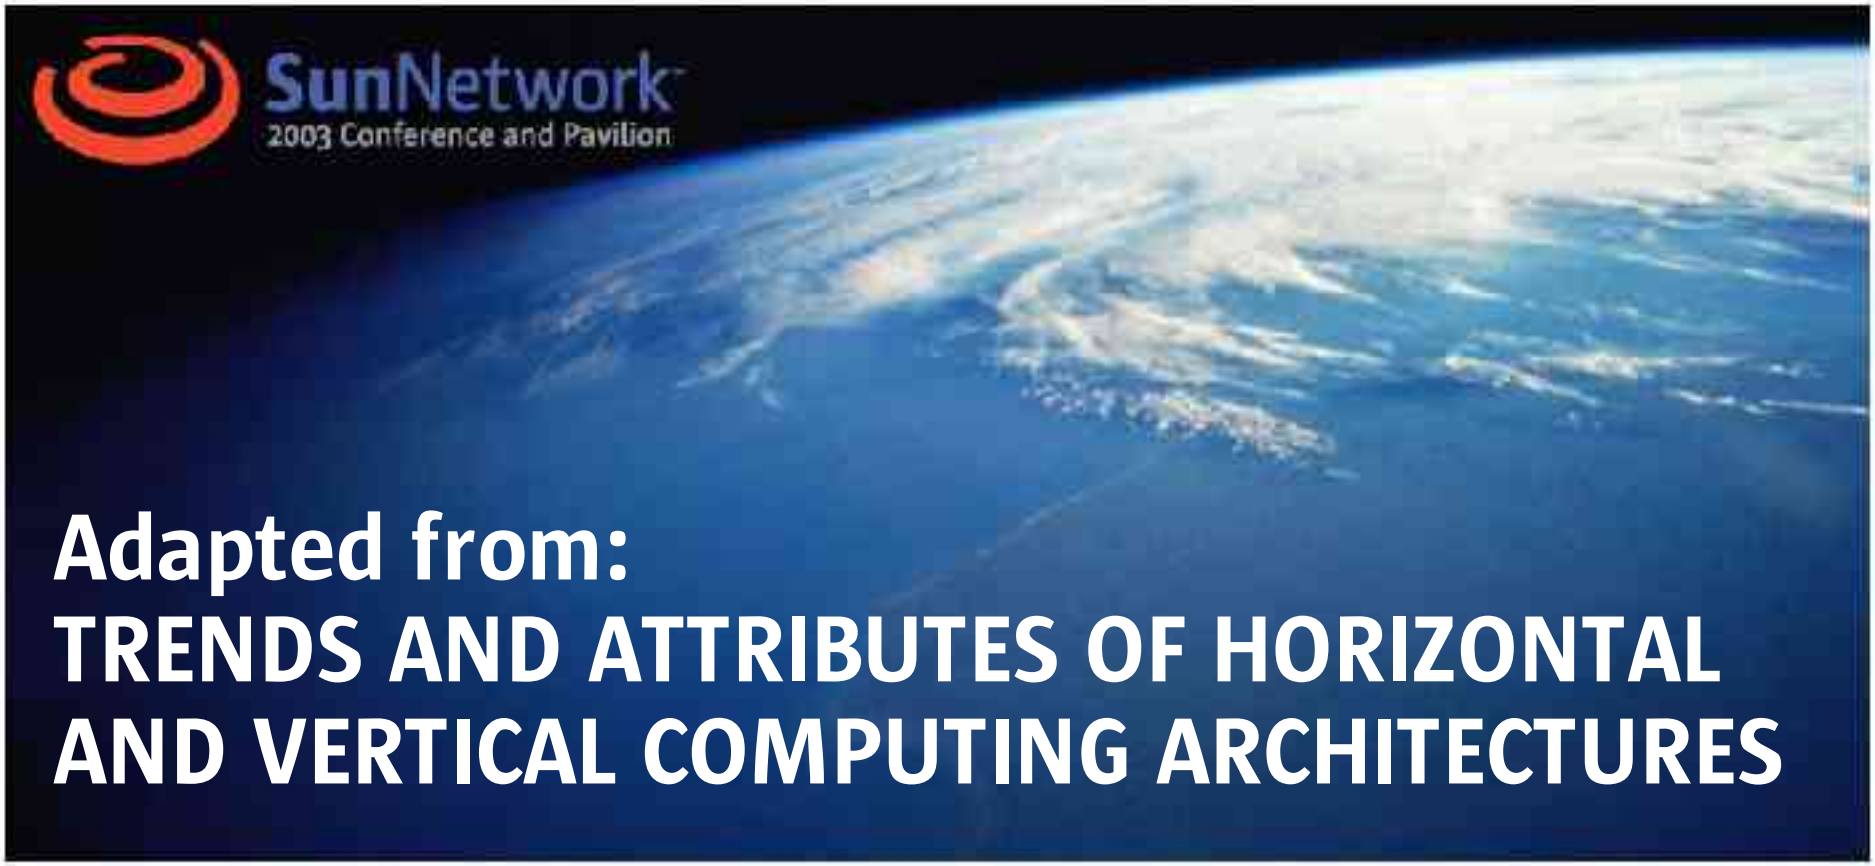

Adapted from:
TRENDS AND ATTRIBUTES OF HORIZONTAL AND VERTICAL COMPUTING ARCHITECTURES

Tom Atwood

Business Development Manager

Sun Microsystems, Inc.

Sun microsystems
We make the network.

Takeaways

- Understand the technical differences between horizontal and vertical architectures
- Understand the difference between cluster for performance and cluster for availability
- Understand the financial impact of cluster databases (e.g., Oracle 9i RAC)

Agenda

- Industry trends
- Architectural definitions
- Performance issues
- TCO issues for database clusters
- Availability issues

N-tier Architecture

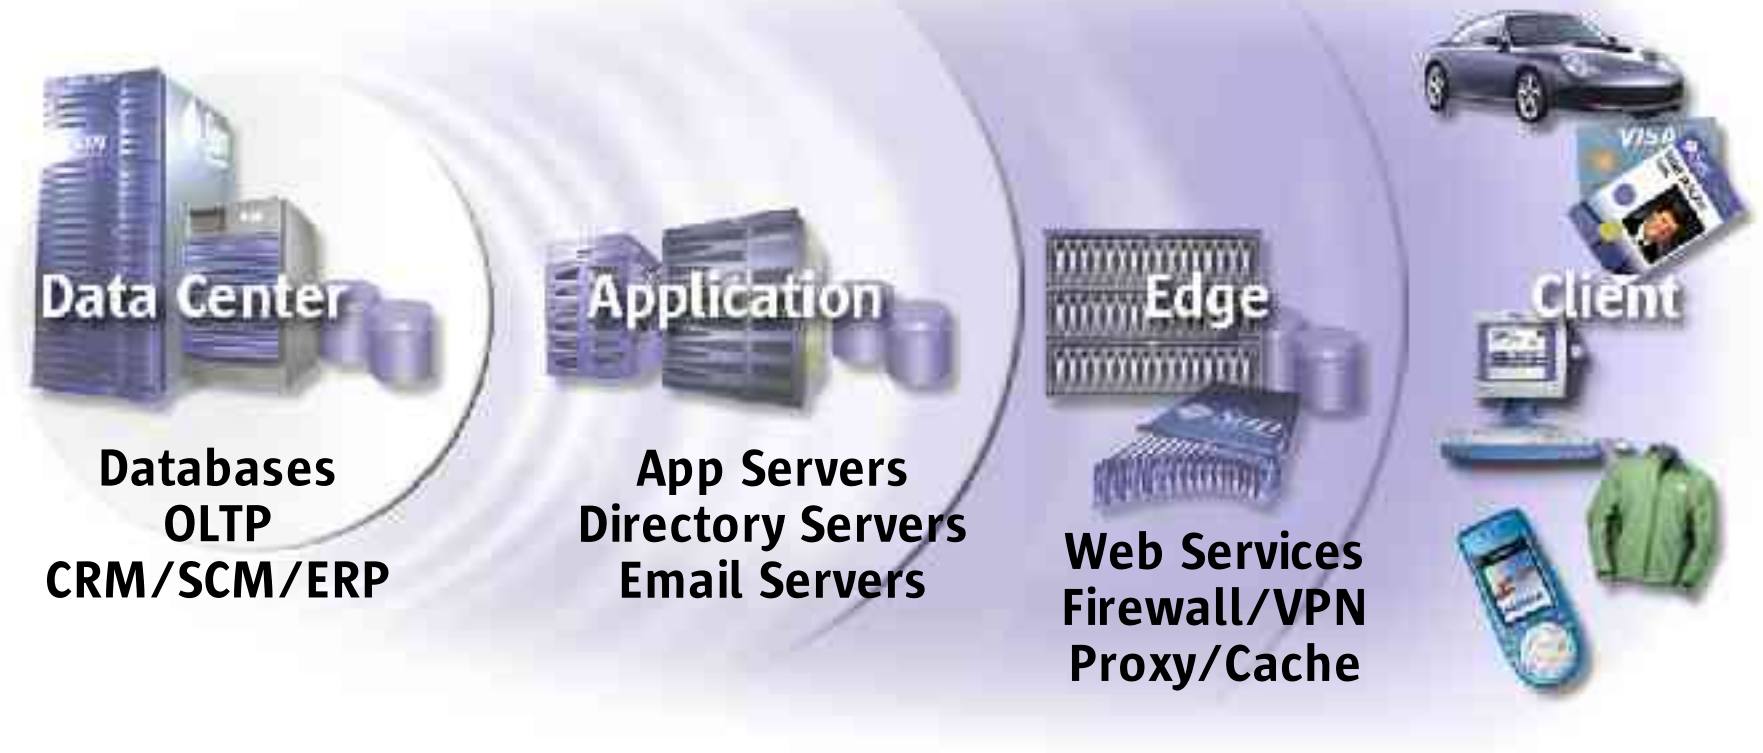

Customer Investment in Vertical and Horizontal Servers

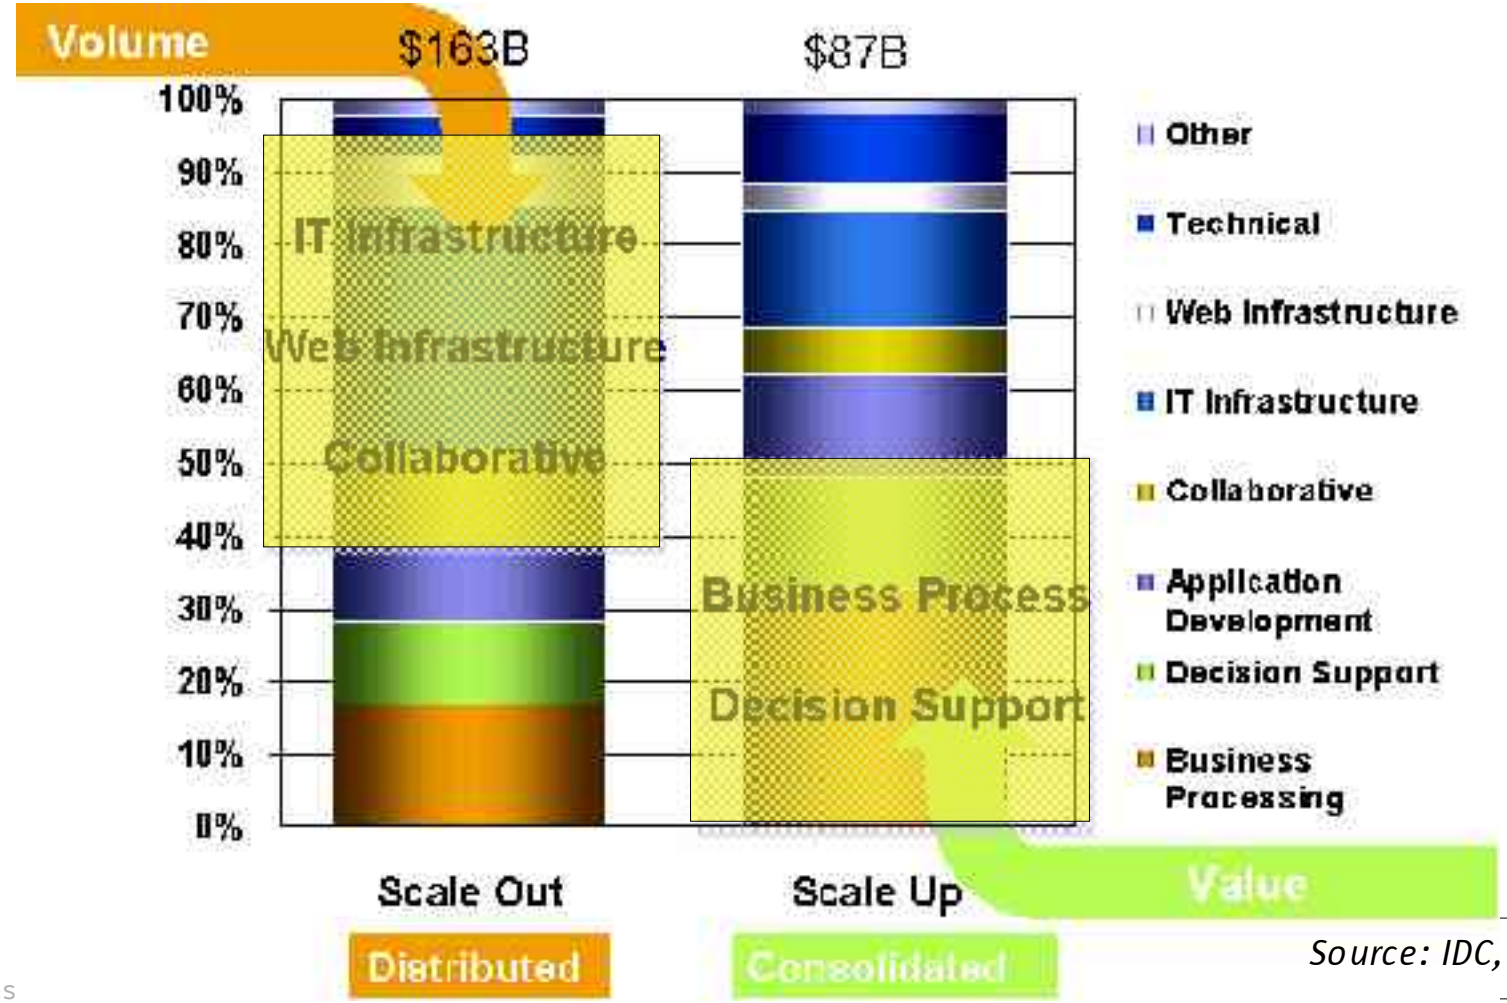

Source: IDC, 2003

Server Usage by CPU Capacity

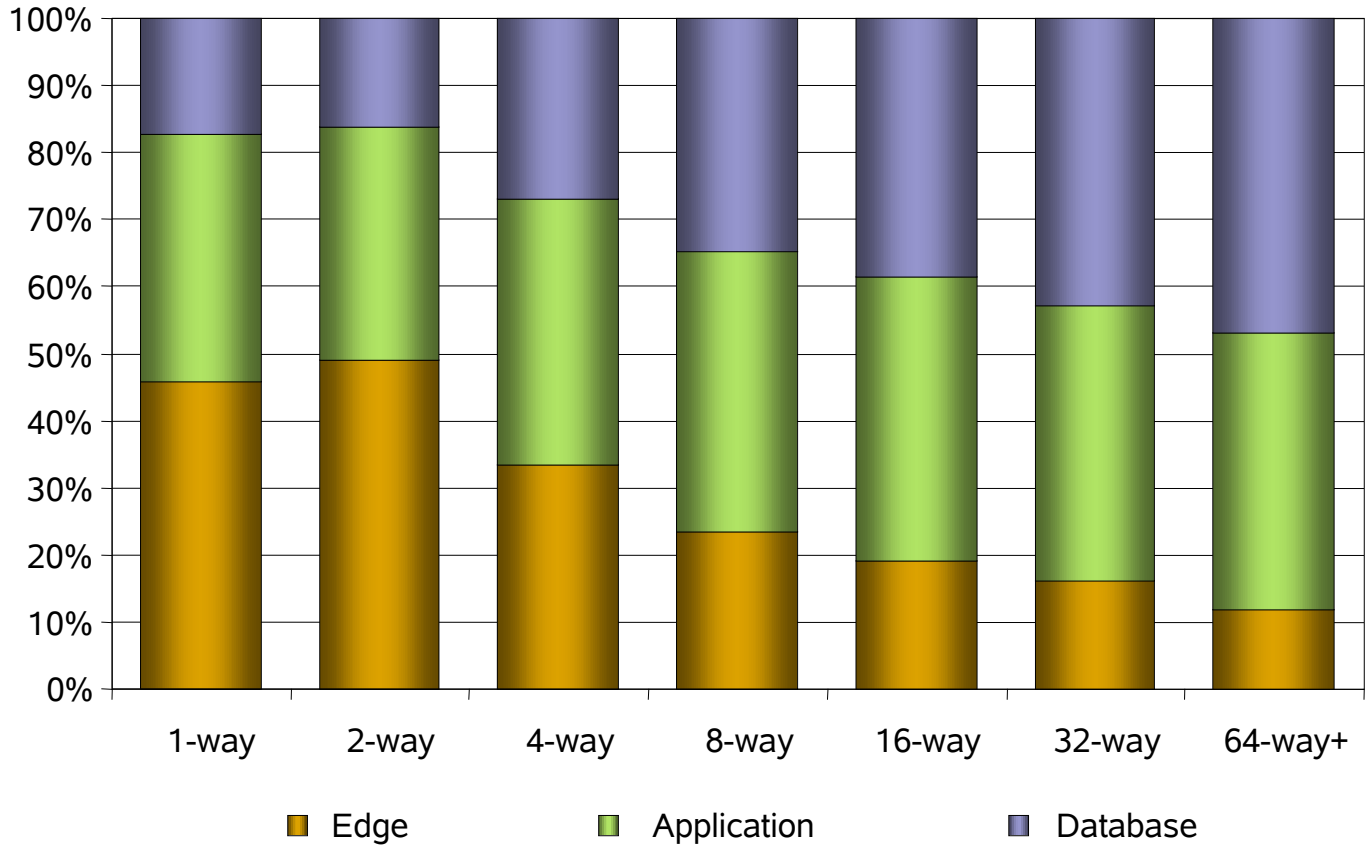

Source: IDC, 2003

SMP Definition

- Symmetric multi-processing
  - Or: shared memory multi-processing
- Shared pool of CPUs
- Single-large memory space
- One copy of the OS

Vertical and Horizontal Attributes

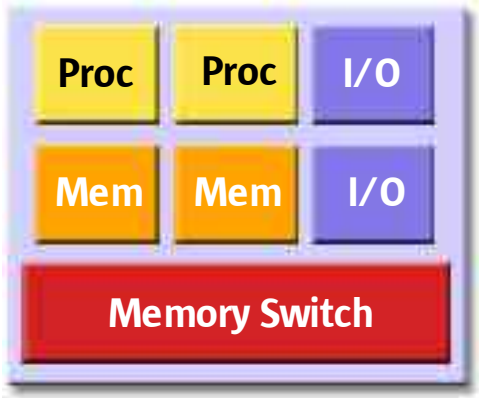

Single OS Instance

Cluster multi-processor

- Loosely coupled
- Standard H/W & S/W
- Highly parallel (web, some HPTC)

Cache-coherent shared-memory multi-processors (SMP)

- Tightly-coupled: high bandwidth, low latency
- Large, workloads: ad-hoc trans proc, data warehousing
- Shared pool processors
- Single large memory

Multiple OS Instances

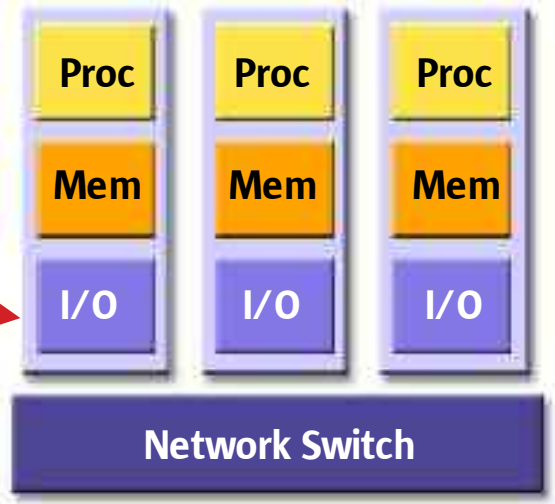

Vertical Scaling

Horizontal Scaling

Vertical vs. Horizontal System Types

- Vertical
  - Large SMPs
  - Clusters of large SMPs
- Horizontal
  - Blades
  - Clusters
  - Large MPPs

Vertical vs. Horizontal Characteristics

- Vertical
  - Large shared memory space
  - Many dependent threads
  - Tightly-coupled internal interconnect
  - High single-system RAS
  - Many standard CPUs
  - Single OS with many CPUs
  - Single-box packaging
  - Many CPUs/floor tile
  - Commodity and proprietary h/w
  - Single-box headroom/growth
  - "In-box" enhancements/upgrades
  - 64-bit
- Horizontal
  - Small non-shared memory space
  - Many independent threads
  - Loosely-coupled external interconnect
  - High RAS via replication
  - Many standard CPUs
  - Many OS's with 1–4 CPUs/OS
  - "Rack and stack" packaging
  - Many cpus/floor tile
  - Commodity hardware
  - Multi-box headroom/growth
  - "New-node" enhancements
  - 32-bit and 64-bit

Vertical vs. Horizontal Applications

- Vertical
  - Large databases
  - Transactional databases
  - Datawarehouses
  - Data mining
  - Application servers
  - HPTC applications (non-partitionable)
- Horizontal
  - Web servers
  - Firewalls
  - Proxy servers
  - Media streaming
  - Directories
  - XML processing
  - JSP applications
  - SSL
  - VPN
  - Application servers
  - HPTC applications (partitionable)

Critical Factors in System Performance

- Processor
  - Capacity and throughput
- System Interconnect
  - Low latency
  - High bandwidth
- Input and output
  - Network and storage, sequential and transactional
- Operating system
  - Required for H/W to scale
- Optimized applications
- System availability

Interconnect Specifications

Network vs. Centerplane/Backplane

- Gigabit ethernet
  - 125MB/sec bandwidth
  - 200,000ns latency
- SCI
  - 200MB/sec bandwidth
  - 4,000ns latency
- Infiniband
  - 250MB/sec-3GB/sec b/w
  - ??

- Sun Fire SMPs
- 9.6GB/sec to 172GB/sec bandwidth
- 200ns to 450ns latency

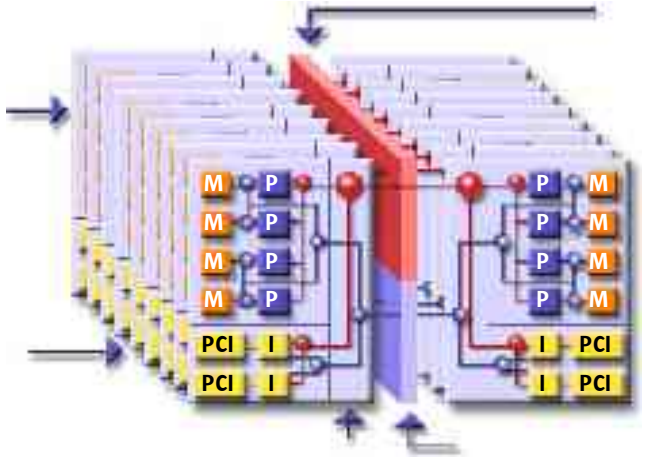

Distributed Memory

Large Workloads Not Handled Effectively

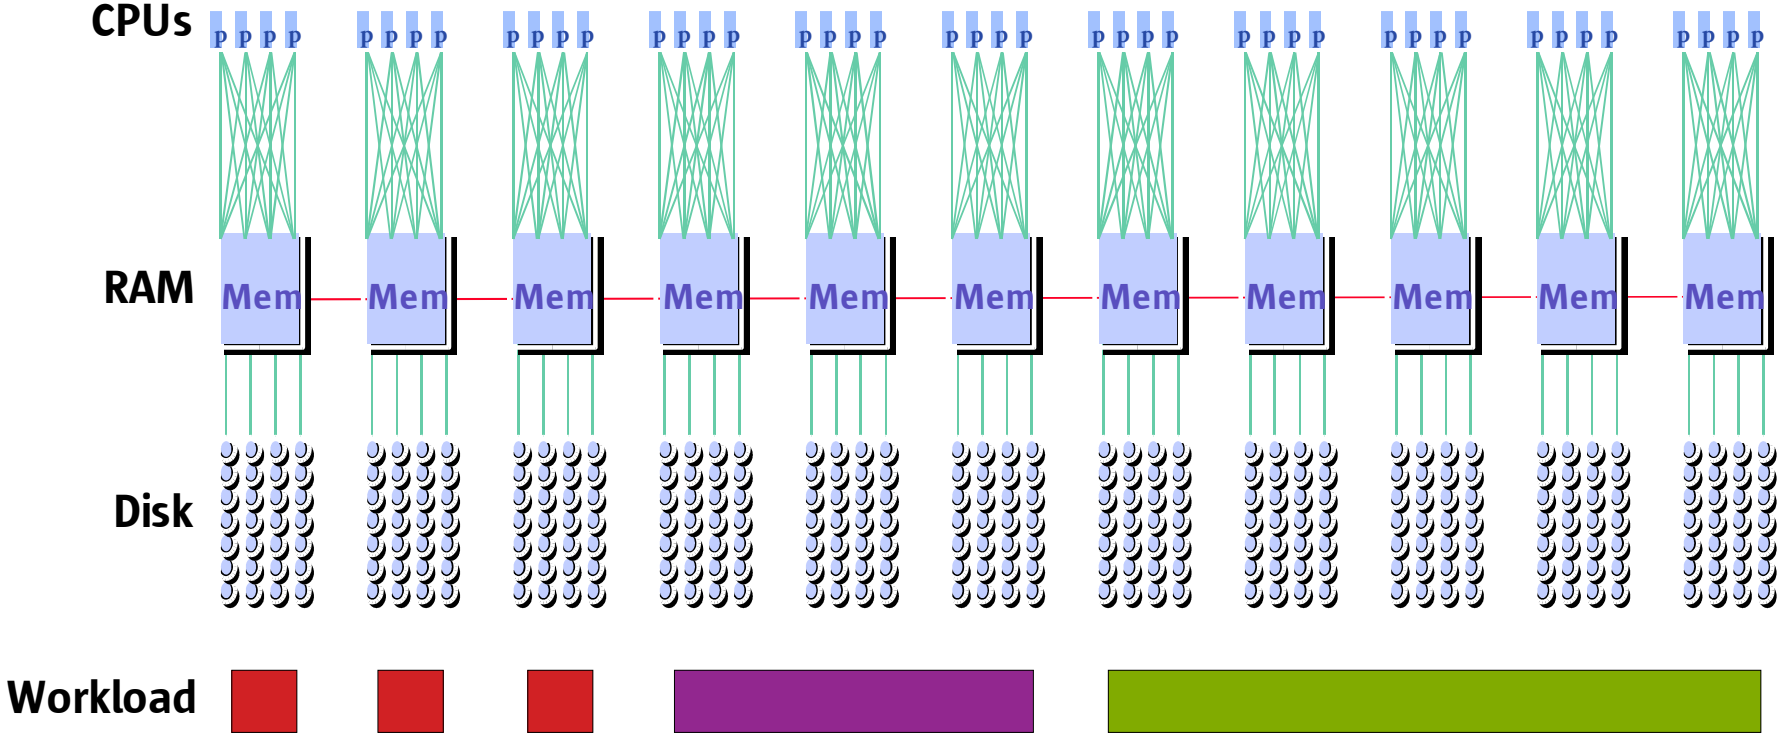

Shared Memory

Large Workloads Handled Effectively

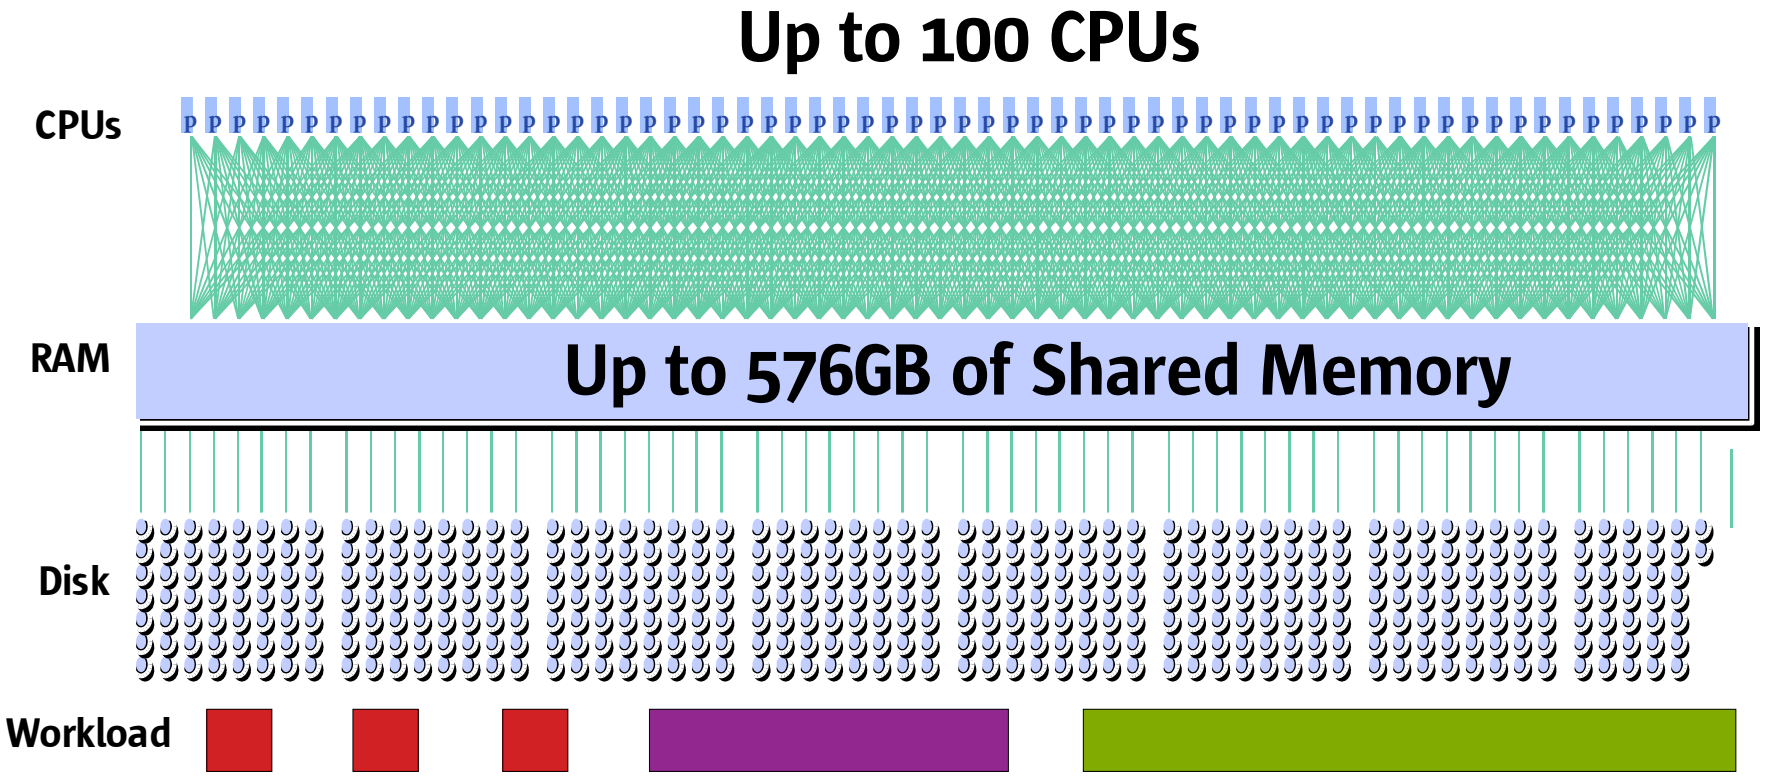

Database Layer Performance

- Horizontal needs to cluster for performance
- Oracle 9i RAC is most common cluster database
  - RAC is “Real Application Clusters”
- Vertical does not need to cluster for performance

What “Scaling” Means

- SMP speedup = x faster than one CPU
- Cluster speedup = x faster than one node
- SMP scaling = speedup/(no. CPUs)
- Cluster scaling = speedup/(no of nodes)
- “Scaling” depends on number of CPUs/nodes used to calculate
- 90% scaling on 2 nodes not the same as 90% scaling on 4 nodes

Scaling Examples

- 1 node to 2 nodes:
  - Scaling factor is .8 or 80%
  - Speedup is 1.8 (1.0, 1.8)
  - Cluster scaling is: 1.8/2 = .9 or 90%
- 4 nodes:
  - Scaling factor is .8 or 80%
  - Speedup is 3.4 (1.0, 1.8, 2.6, 3.4)
  - Cluster scaling is: 3.4/4 = .85 or 85%

Scaling of Clusters

- 90% Scaling on 2-way Cluster of 4-way servers gives 23x Speedup at 52 CPUs (Oracle 91 RAC)*
- 95% Scaling gives 23x speedup on 24 CPU SMP (Oracle)**
- Speedup is how much faster than one CPU

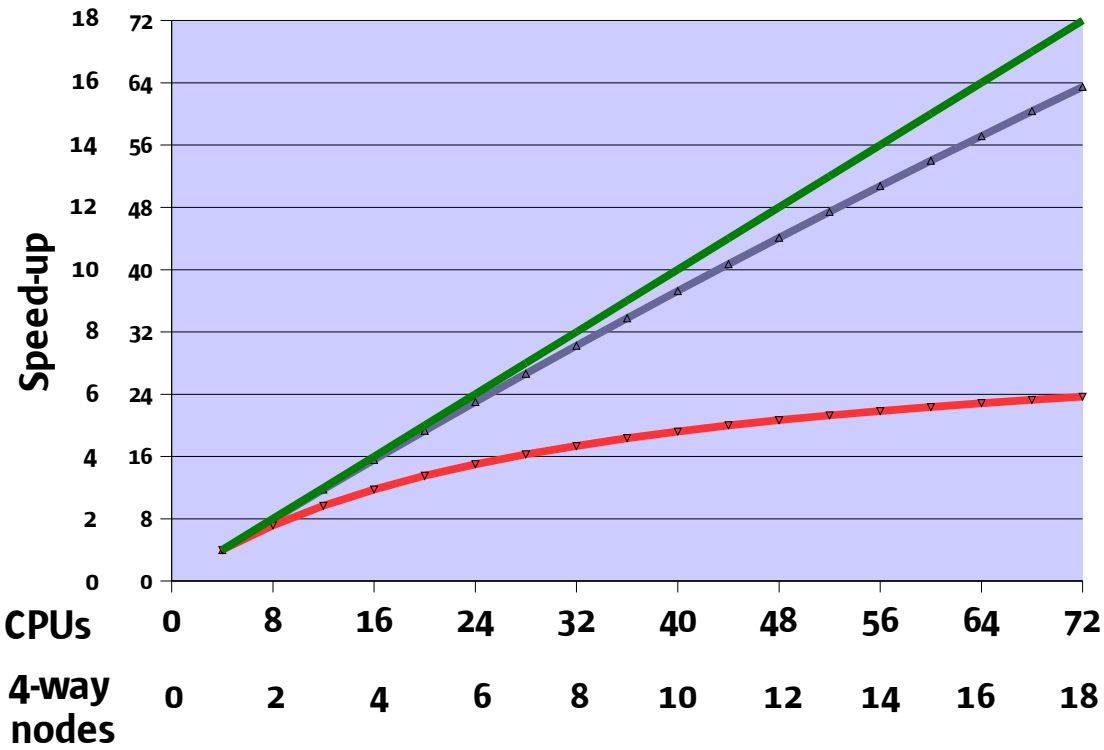

Linear scaling

95% SMP scaling (.95=23/24CPUs)

90% cluster scaling (.90=1.8/2-servers)

* 90% scaling from two benchmarks:
www.sap.com/benchmark/others.htm#Round%Robin
www.oracle.com/ip/deploy/database/theme_pages/index.html?rac_01242003.html

** 95% scaling from Internal Sun benchmarks

Application Layer Performance

- Application layer is normally stateless
  - Data resides in database layer
  - Consolidates connections to the database
- More processors needed than DB layer
  - SAP R/3 uses about 10 app CPUs per each DB CPU
  - Oracle Apps uses about 5 app CPUs per each DB CPU
- Scaling is not an issue
  - Replication of instances meets performance requirements
- Acquisition costs not affected by different software licenses

SAP App Layer Example

Fewer Large App Servers (%Use)

Server Consolidation:

- Higher initial costs
- Less Complexity
- Fewer Boxes
- Possible app contention
- Possibly lower admin costs
  Possibly fewer licenses,...

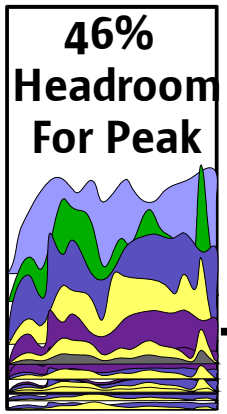

DB Server

Many Little App Servers (%Use)

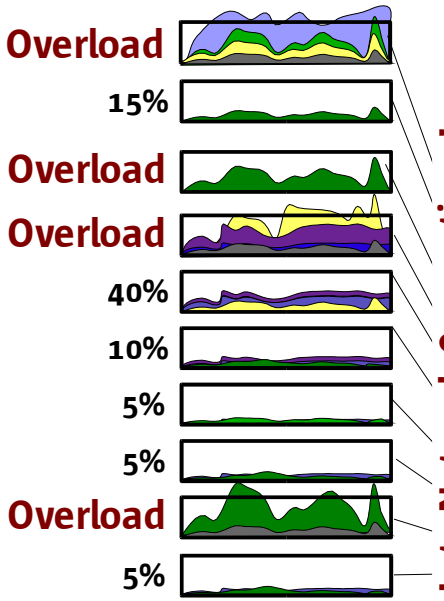

Server Replication:

- Lower initial costs
- "Hot spots"
- Less predictable
- Possibly higher management costs

DB Server

- Customer Evaluation: SAP Application Servers
- 50% Less HW 24-way vs. 2-ways
- Load Balancing on SAP Sessions at Login — Static
- Larger servers had more headroom for peak workloads
- Batch Jobs (Payroll, MRP...) can't span several app servers

Application Layer Management

- If database requires 20 processors
  - Oracle Apps application layer: 100 processors
  - SAP R3 application layer: 200 processors
- 50–100 2-way servers or 5–10 20-way servers?
- More OS instances may increase costs
- More application instances may increase costs
- Horizontal solution may have lower TCA but higher TCO

Consolidated Solutions

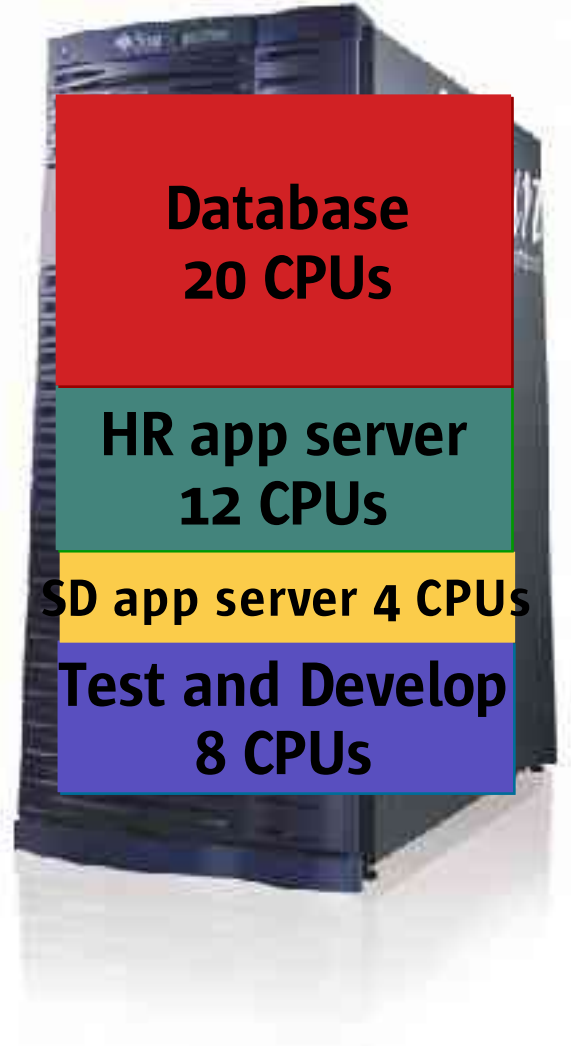

Presentation Layer Performance

- Applications rarely scale
  - Applications cannot use multiple CPUs
- Applications are not data intensive
  - Commonly stateless
- Internode communication minimal
- Cluster costs are not an issue
- Horizontal has lower acquisition cost
- Horizontal has sufficient performance

How to Lower TCO

- Lower acquisition costs
  - Use commodity components
- Lower complexity
  - Fewer OS “flavors”
  - Fewer OS instances
- Better utilize resources
  - Virtualize resources
  - Automate resource migration
- Centralize
  - Data centers
  - Disaster recovery
  - Backup

Application Lifecycle TCO

- TCO must be managed over entire application lifecycle
- Less than 20% is hardware acquisition

TCO Implications of Cluster Scalability

- Oracle 9i RAC needs more CPUs than Oracle 9i
- Oracle 9i RAC per CPU pricing higher than Oracle 9i
- Acquisition costs (h/w and s/w) may be higher for horizontal scaling
- E.g., Need 13 x 4-CPU servers to equal 24-way SMP server
  - 52 CPUs vs. 24 CPUs

Availability Continuum

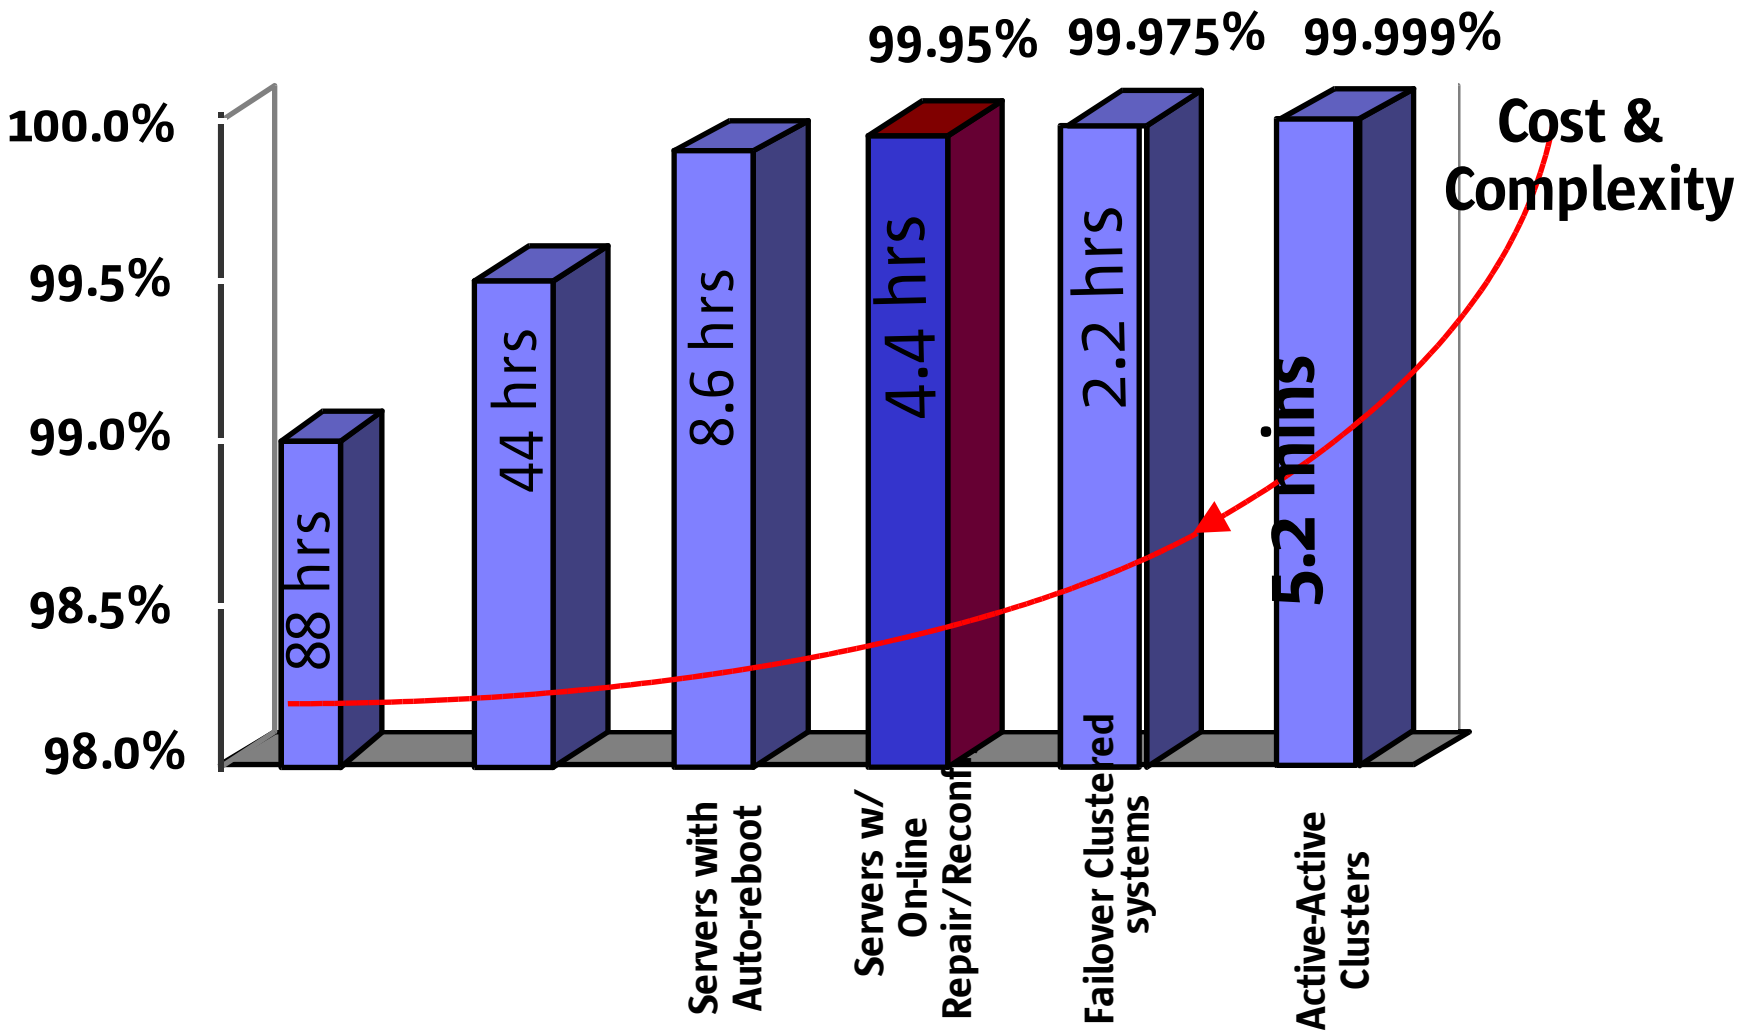

Availability Strategies

Different for Horizontal and Vertical

- Large SMPs: High single system RAS
  - Full hardware redundancy
  - Online serviceability
  - Error checking and correction
  - HA failover (standby server)
  - Cluster for highest availability requirements
- Horizontal nodes: Clusters and Replication
  - Cluster for availability for databases
  - Replicate workload for non-DB applications

TCO Implications of Availability

Level of Availability Affects TCO

- 99.95% single SMP may be sufficient
- Greater than 99.95% requires cluster
  - Standby failover (HA) may be sufficient
    - Database active on only one node, other standby
    - Database needs to start up
    - For Oracle does not require RAC licenses
  - Highest avail. level requires active cluster
    - 2 or more nodes active
    - Fast failover
    - For Oracle RAC licenses required

Large vs. Small Clusters

Hardware and Software Acquisition Costs

- 1 x 6800 20-way server: $826,360
  - Oracle 9i licenses
- 2 x 6800 12-way servers: $1,461,360
  - Same performance as 1 x 20-way 6800
  - With Oracle 9i RAC licenses
- 8 x 480 4-way servers: $1,465,600
  - Same performance as 1 x 20-way 6800
  - With Oracle 9i RAC licenses

RAC Calculator Parameters: 90% scaling, 10% decay, 50% Oracle discount, 40% 6800 discount, 20% 480 discount

Clustering for Scalability vs. Clustering for Availability

IDC Survey Data:

Source: IDC, 2003

"Even customers most familiar with clustering software and clustered systems continue to be very cautious regarding the replacement of large scale-up systems with clusters of scale-out systems."

— Matthew Eastwood, IDC

Availability Summary

- For highest availability: (<99.975)
  - Horizontal or vertical clustering
  - Two large-node clusters similar cost as multiple small node clusters
- For lower availability
  - Vertical with HA failover will cost less
  - Single vertical node may be sufficient

Vertical and Horizontal Summary

- Vertical scaling
- Big memory
- Big disk I/O bandwidth
- Big RAS
- Protect state
- SAN

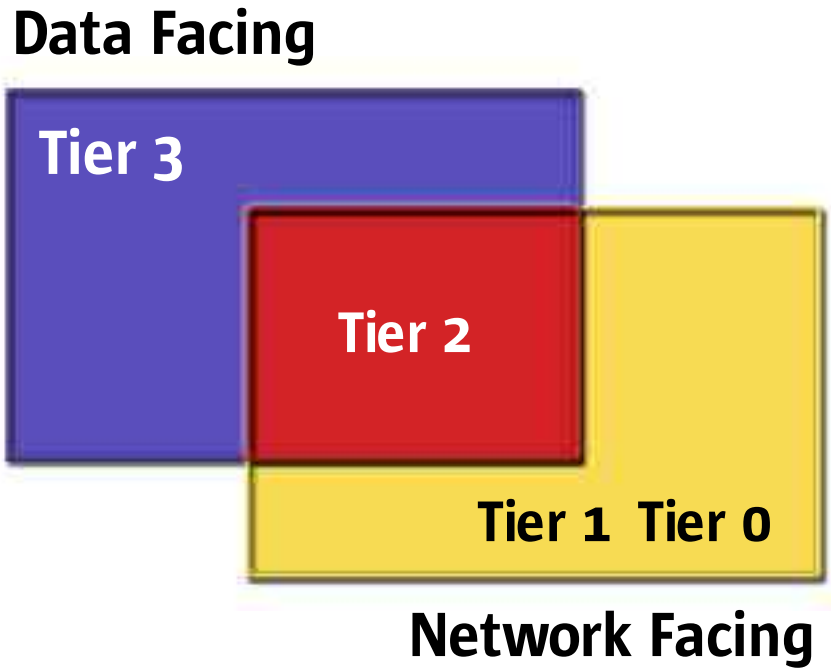

- Horizontal scaling
- Med memory
- Big network I/O bandwidth
- Replication
- Stateless
- NAS

Horizontal vs. Vertical Summary

- Horizontal ideal for web-tier
  - Performance and acquisition cost
- Vertical well-suited for database tier
  - Performance and acquisition cost
- Both can be used for application tier
- Clustering good for availability but not for performance